

An Introduction to ARCH/GARCH Modeling Using the AUTOREG Procedure

Donna E. Woodward
SAS Institute

ABSTRACT

Many economic and financial time series exhibit a phenomenon known as heteroscedasticity, in which the variance of the series changes over time. The changes in variability are often based on recent past shocks, which cause large (small) fluctuations to cluster together. The recognition that this phenomenon could be used to more accurately define the forecast error variance led to the introduction of the ARCH or Autoregressive Conditional Heteroscedasticity model by Engle in 1982. Further refinements to this model resulted in the class of models known as GARCH or Generalized ARCH models. This paper takes a brief look at the history of these models and illustrates how they can be estimated and forecast using the SAS® System.

INTRODUCTION

Although researchers have known for many years that the variance of economic and financial data changes over time, Engle was the first to formulate a model that captured the essence of this heteroscedasticity. The popularity of this model can be seen by the voluminous attention it has received in the literature and the various extensions that have been made to it since its introduction slightly more than a decade ago.

Most traditional time series models, such as ARIMA models, emphasize modeling and forecasting the conditional mean of a series, assuming that the conditional variance is constant. In other words, for a simple AR(1) model, they assume that the forecast for $y_t = E[y_t | y_{t-1}]$, and the conditional variance is simply σ^2 . Engle (1982) argued that more accurate forecast intervals could be obtained by using the recent past to affect not only the conditional mean, but also the conditional variance. He called this class of models AutoRegressive Conditional Heteroscedasticity or ARCH models because the form of the heteroscedasticity is conditional on previous information. Figure 1 exhibits this behavior for the mean-deleted inflation rate data based on the U.S. CPI from January, 1970 through February, 1995.




Figure 1. U.S. Monthly Inflation Rate

This paper will describe the ARCH and GARCH models and will illustrate how they can be estimated with PROC AUTOREG in the SAS/ETS® product. It will first present these models based on generated data and will then conclude with a complete analysis of the above CPI data. It will also describe how to obtain predicted values of not only the conditional mean but also of the conditional error variance. Finally, a brief summary of several new GARCH modeling features forthcoming in the 6.11 release of the SAS System will be provided.

ARCH MODELS

The ARCH model can be presented in a variety of contexts. The simplest case is when the model is applied directly to the past realizations of the series. Under the assumption of normality, Engle (1982) presents this model as:

$$y_t = \varepsilon_t \sqrt{h_t}$$

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2$$

where $y_t | \Psi_{t-1} \sim N(0, h_t)$, and

Ψ_{t-1} is the information set at time $t-1$.

This model can easily be extended to the case of a dynamic regression model where y_t is a function of exogenous and possibly lagged endogenous variables. The ARCH process is then applied to the errors of this regression model, which is written as:

$$y_t = x_t' \beta + \varepsilon_t,$$

$$\varepsilon_t = e_t \sqrt{h_t}$$

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2$$

where $y_t | \Psi_{t-1} \sim N(x_t' \beta, h_t)$, and $e_t \sim IN(0,1)$.

To ensure that the conditional error variance, h_t , is positive, the restrictions that $\omega > 0$ and $\alpha_i \geq 0, i=1,2,\dots,q$, are imposed. Engle shows that under these non-negativity constraints and when $\sum_{i=1}^q \alpha_i < 1$, the ARCH(q) model has positive and finite unconditional variance $V(\varepsilon_t) = \omega / (1 - \sum_{i=1}^q \alpha_i)$.

From this model, we can see that large disturbances cause the conditional variance, h_t , to increase. The sign of the disturbance is unimportant because we are considering only the squared errors. The length of time the shock effects the conditional variance depends on the value of q . Large values of q cause a shock to have a longer impact on the conditional variance, resulting in wider forecast intervals during those affected periods.

EXAMPLE 1

Bera and Higgins (1992) provide a nice comparison of ARCH processes by generating a Gaussian white noise series, $e_t \sim N(0,1)$, and multiplying it by two different conditional standard deviations, $\sqrt{h_t}$, to illustrate an ARCH(1) and an ARCH(4) process. Following is sample SAS code that shows how to generate data from these processes and then plot them using the SAS/GRAPH® GPROT procedure:

ARCH(1):

$$h_t = .1 + .9\varepsilon_{t-1}^2$$

ARCH(4):

$$h_t = .1 + .36\varepsilon_{t-1}^2 + .27\varepsilon_{t-2}^2 + .18\varepsilon_{t-3}^2 + .09\varepsilon_{t-4}^2$$

```
data gen1;
  retain l1arch1 0
         l1arch4 l2arch4 l3arch4 l4arch4 0;
  do time = -20 to 500;
    e=rannor(12345);
    ht1 = .1+.9*l1arch1;
    ht4 = .1+.36*l1arch4+.27*l2arch4+
          .18*l3arch4+.09*l4arch4;
    arch1=e*sqrt(ht1);
    arch4=e*sqrt(ht4);
    if time > 0 then output;
    l1arch1=arch1*arch1;
    l1arch4=l3arch4; l3arch4=l2arch4;
    l2arch4=l1arch4; l1arch4=arch4*arch4;
  end;
run;

symbol1 v=none i=join;
proc gplot data=gen1;
  axis1 order=(-7 to 7);
  plot (e arch1 arch4)*time/vaxis=axis1;
run;
```

Figure 2. Generation and Plotting of White Noise, ARCH(1) and ARCH(4) data.




Figure 3. White Noise ~ N(0,1) process




Figure 4. ARCH(1) process




Figure 5. ARCH(4) process

Notice the difference in the variability in the three plots--the white noise process, as expected, has no apparent pattern that can be modeled. However, the ARCH(1) and ARCH(4) series have clusters of large (small) variance. The effect of a shock on the variability lasts longer for the ARCH(4) process than for the ARCH(1) process. This changing variability is what ARCH models attempt to explain.

GARCH MODELS

Since Engle's introduction of the ARCH model, many extensions of this model have been formulated. The most commonly cited one is the Generalized ARCH or GARCH model proposed by Bollerslev (1986). Because the usual ARCH model often requires large values of q to adequately account for the heteroscedasticity, a lag structure must often be imposed on the parameters in the conditional variance in order to meet the non-negativity constraints (Engle, 1982).

Bollerslev proposed an extension of the ARCH model to the GARCH model that resembles the extension of a MA model to an ARMA model. It allows for a more parsimonious model specification without having to impose a somewhat arbitrary lag structure on the α_i 's. In the ARCH(q) process, the conditional variance is solely a function of the past sample variances. In the GARCH(p,q) process, however, the conditional variance is defined in terms of both lagged sample variances and past values of the conditional variance (Bollerslev, 1986). As in the case of the ARCH model, the GARCH process

may be applied to either the past realizations of the series or to the errors of a dynamic regression model.

The conditional variance for the GARCH(p,q) model is:

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \gamma_j h_{t-j}$$

where $p \geq 0, q > 0, \omega > 0, \alpha_i \geq 0, \gamma_j \geq 0$.

Tables 1 and 2 illustrate how the extension of ARCH to GARCH parallels the extension of the MA model to the more general ARMA model. As in ARMA models, under certain conditions, an infinite order MA process can be written as a low order AR process; so it follows that a high order ARCH process can, under certain conditions (Bollerslev, 1986), be written as a low order GARCH model.

$$y_t = \mu + a_t + \sum_{i=1}^q \Theta_i a_{t-i}$$

$$y_t = \mu + a_t + \sum_{i=1}^q \Theta_i a_{t-i} - \sum_{j=1}^p \Phi_j y_{t-j}$$

Table 1. Extension of MA model to ARMA model

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2$$

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \gamma_j h_{t-j}$$

Table 2. Extension of ARCH model to GARCH model

The most commonly used and easiest to specify GARCH model is the GARCH(1,1) model. Its conditional variance is expressed as:

$$h_t = \omega + \alpha_1 \varepsilon_{t-1}^2 + \gamma_1 h_{t-1}$$

where $\omega > 0, \alpha_1 \geq 0, \gamma_1 \geq 0$.

EXAMPLE 2

The following code in Figure 6 generates and plots a data set with the GARCH(1,1) conditional variance given by:

$$h_t = 1 + 2\varepsilon_{t-1}^2 + 0.6h_{t-1}$$

Figure 7 shows the plot of this generated data.

```

data gen2;
  retain llesq l1h 0;
  do time = -20 to 500;
    e=rannor(12345);
    ht = .1+.2*llesq + .6*l1h;
    garch1_1=e*sqrt(ht);
    if time > 0 then output;
    llesq=garch1_1*garch1_1; l1h=ht;
  end;
run;

proc gplot data=gen2;
  axis1 order=(-7 to 7);
  plot garch1_1*time/vaxis=axis1;
run;

```

Figure 6. Generation and Plotting of GARCH(1,1) data




Figure 7. GARCH(1,1) process

DETERMINING THE TYPE AND ORDER OF THE ARCH/GARCH PROCESS

Weiss (1984) has shown the importance of detecting the presence of ARCH disturbances in times series data. Specifically, he showed that ignoring the presence of heteroscedasticity not only causes the OLS estimates to be inefficient, but it also could result in an overparameterized ARMA model. With this in mind, we need to address the problem of detecting the existence of ARCH disturbances and determining the order of the process.

In his seminal paper, Engle (1982) proposed the use of the Lagrange multiplier (LM) test to ascertain the presence of ARCH. The null hypothesis is: $H_0: \alpha_1 = \alpha_2 = \dots = \alpha_q = 0$, v.s. the alternative,

$H_1: \text{ARCH}(q)$. Lagrange multiplier test statistics are computed for various values of q and are formed by obtaining the OLS residuals, squaring them, and then regressing them on a constant and q lags of the squared residuals. The test statistic for lag q is approximated by $T \cdot R^2$, where T is the number of residuals and R^2 is the coefficient of multiple determination for the current lag model. $T \cdot R^2$ follows a X_q^2 distribution.

While the LM test allows us to test for the presence of heteroscedasticity, it does not provide much information as to the form of the underlying process (Bera and Higgins, 1992). However, it is often true that if the LM statistic is significant for higher order lags, then it may indicate that a GARCH model could more parsimoniously predict the conditional error variance than a higher order ARCH model.

Another statistic, the Portmanteau Q-statistic, used to test for nonlinear effects in time series models (McLeod and Li, 1983) may also be used to detect the presence of ARCH. Bollerslev (1986) showed that the GARCH(p,q) model could be written as:

$$\varepsilon_t^2 = \omega + \sum_{i=1}^n (\alpha_i + \gamma_i) \varepsilon_{t-i}^2 - \sum_{j=1}^p \gamma_j \eta_{t-j} + \eta_t$$

where $\eta_t = \varepsilon_t^2 - h_t$, and $n = \max(p, q)$.

This representation of the GARCH model follows the standard ARMA form for the squared series, therefore, conventional methods used to identify ARMA processes may be used to determine the presence of ARCH. This includes using the sample autocorrelation and partial autocorrelation functions, as well as the Portmanteau Q-test. The Q-test is based on the sample autocorrelation function (SACF) of the squared OLS residuals for various lags, and follows a X_q^2 distribution. The null hypothesis is, again, white noise.

Although the Lagrange multiplier and Portmanteau Q-statistics are helpful in detecting ARCH, it is still often difficult in practice to determine the order of the process. One final method for determining the order of the model is simply to fit several competing models and then compare the AIC and SBC values for these competing models. The AIC (Akaike Information Criterion) and SBC (Schwarz Bayesian Criterion) are not meaningful in and of themselves, but are used solely when comparing models. They are relative measures of model adequacy, similar in principle to the SSE. For a given set of data, the model yielding the smallest AIC and/or SBC is the best fitting model among those estimated.

EXAMPLE 3

To illustrate the difficulties in determining the presence and order of the ARCH/GARCH process, let us examine the LM and Q statistics computed by the ARCHTEST option in the AUTOREG procedure for the data generated by the ARCH(1) process. After examining the results of the ARCHTEST option, we will fit several competing models and then select the best one based upon the AIC and SBC. Figure 8 contains the PROC AUTOREG code for obtaining the LM and Q statistics. Note that the NOINT option is used to suppress the automatic fitting of an intercept term, because we know the generated ARCH1 series is zero mean. Figure 9 contains the resulting listing file.

```
proc autoreg data=gen1;
  model arch1=/archtest noint;
run;
```

Figure 8. PROC AUTOREG and the ARCHTEST option

The SAS System					
Autoreg Procedure					
Dependent Variable = ARCH1					
Ordinary Least Squares Estimates					
SSE	386.8793	DFE	500		
MSE	0.773759	Root MSE	0.879635		
SBC	1290.691	AIC	1290.691		
Reg Rsq	0.0000	Total Rsq	0.0000		
Durbin-Watson	1.8521				
NOTE: No intercept term is used. R-squares are redefined.					
Q and LM Tests for ARCH Disturbances					
Order	Q	Prob>Q	LM	Prob>LM	
1	123.327	0.0001	122.627	0.0001	
2	194.139	0.0001	133.772	0.0001	
3	202.603	0.0001	141.914	0.0001	
4	203.743	0.0001	142.262	0.0001	
5	203.879	0.0001	142.935	0.0001	
6	204.095	0.0001	143.162	0.0001	
7	204.962	0.0001	143.374	0.0001	
8	210.529	0.0001	146.304	0.0001	
9	235.063	0.0001	157.619	0.0001	
10	322.349	0.0001	191.324	0.0001	
11	352.020	0.0001	201.427	0.0001	
12	373.869	0.0001	201.922	0.0001	

Figure 9. LM and Q-statistics for ARCH(1) series

Figure 9 shows that all 12 lags used in computing the Lagrange multiplier and Portmanteau Q-statistics are significant. This implies that a higher order ARCH or a GARCH(1,1) model might be a good choice to fit. Since we know this data was generated by an ARCH(1) process, we will fit and compare two models using the AUTOREG procedure: a GARCH(1,1) model and an ARCH(1) model. Figure 10 provides the SAS code for fitting these two models.

```
proc autoreg data=gen1;
  model arch1=/garch=(p=1,q=1) noint;
  model arch1=/garch=(q=1) noint;
run;
```

Figure 10. Fitting a GARCH(1,1) and an ARCH(1) model using PROC AUTOREG

An ARCH or GARCH model is fit in the AUTOREG procedure by specifying the GARCH= option on the MODEL statement, followed by either the $q=$ option or the $p=$ and $q=$ options. Maximum likelihood estimation is automatically used whenever an ARCH/GARCH model is fit. The listing file for the GARCH(1,1) model, suggested by the LM and Q-statistics, is shown below in Figure 11. The OLS results have been omitted from this figure since they are identical to what was shown in Figure 9.

Dependent Variable = ARCH1					
GARCH Estimates					
SSE	386.8793	OBS	500		
MSE	0.773759	UVAR	0.671468		
Log L	-357.561	Total Rsq	0.0000		
SBC	733.7656	AIC	721.1218		
Normality Test	8.5334	Prob>Chi-Sq	0.0140		
Variable	DF	B Value	Std Error	t Ratio	Approx Pr
ARCH0	1	0.085511	0.0167	5.134	0.0001
ARCH1	1	0.830992	0.0972	8.550	0.0001
GARCH1	1	0.041658	0.0544	0.766	0.4434

Figure 11. GARCH(1,1) parameter estimates

PROC AUTOREG first provides several summary statistics, including the sum of squared errors (SSE), the number of observations included in the analysis (OBS), the mean squared error (MSE), the unconditional variance (UVAR), the value of the log likelihood function (Log L), the Total R-square for the model, the SBC and AIC, and a test for normality and its associated p-value. Most of these summary statistics are self-explanatory, but two will be discussed in more detail: the unconditional variance and the Normality test.

The unconditional variance (UVAR) for the stationary GARCH(1,1) model is computed as:

$V(\varepsilon_t) = \omega / (1 - \alpha_1 - \gamma_1)$, therefore, it is based on the GARCH parameter estimates. In this case, it is computed as: $.085511 / (1 - .830992 - .041658) = .671468$.

The Normality test (Bera and Jarque, 1982), also presented in the GARCH Estimates table, can be used to detect departures from the assumption of normality, which could indicate misspecification of the error distribution. It is based on the discounted residuals,

$\varepsilon_t / \sqrt{h_t}$. In this case, the null hypothesis of normality is rejected, because the associated p-value is small ($p=.014$). This could indicate that either the model is not correct or that the data do not come from a Normal distribution. Note that this test is only asymptotically valid, therefore, it may reject more often for small sample sizes.

The final table shown in Figure 11 is the table of parameter estimates and their standard errors, t-statistics and p-values. The p-value is associated with the test that the parameter estimate is equal to 0. In this output, ARCH0 is the parameter, ω , ARCH1 is the parameter, α_1 , and GARCH1 is the parameter, γ_1 . The ARCH0 and ARCH1 parameter estimates are both significant, however, the GARCH1 parameter estimate is not ($p=.4434$). Figure 12 shows the output from fitting the simple ARCH1 model.

Dependent Variable = ARCH1					
GARCH Estimates					
SSE	386.8793	OBS	500		
MSE	0.773759	UVAR	0.613995		
Log L	-357.761	Total Rsq	0.0000		
SBC	727.9512	AIC	719.522		
Normality Test	9.4268	Prob>Chi-Sq	0.0090		
Variable	DF	B Value	Std Error	t Ratio	Approx Pr
ARCH0	1	0.095074	0.0108	8.809	0.0001
ARCH1	1	0.845155	0.0964	8.765	0.0001

Figure 12. ARCH(1) parameter estimates

The ARCH0 and ARCH1 parameter estimates are both significant for this model, however, we also reject the null hypothesis of normality of the discounted residuals for the ARCH(1) model. Because this is generated data, we know that the series follows a Normal distribution, whose conditional variance was generated by the process:

$h_t = 1 + .9\varepsilon_{t-1}^2$. The parameter estimates are relatively close to the parameter values used to generate the data. Note that when the sample size was increased to 1000, the Normality test was no longer rejected.

Based on the tests of significance of the parameter estimates and the AIC/SBC values from these two competing models, the ARCH(1) model fits the data better.

ANALYSIS OF U.S. MONTHLY INFLATION RATE DATA BASED ON THE C.P.I.

The data plotted in Figure 1 reflects the mean-corrected monthly inflation rate for the United States from January, 1970 to February, 1995. It is derived from the monthly CPI by performing the operations shown in Figure 13. The residual variable, INFLATN, output by the AUTOREG procedure, is the variable we will analyze. Not all steps in the analysis will be shown, however, the entire process will be described.

```
filename dat '<pathname>';

data cpi(keep=date dflogcpi);
  infile dat;
  input @1 year 4. @6 month 2. cpi;
  date=mdy(month,1,year);
  dflogcpi=dif(log(cpi));
  format date date7.;
run;

/* obtain mean-deleted series */
proc autoreg data=cpi noprint;
  model dflogcpi=;
  output out=meandel r=inflatn;
run;
```

Figure 13. Input and transform CPI data

The ARIMA procedure in the SAS/ETS product was run on the data to aid in identifying any autoregressive or moving average patterns in the data. Based on the resulting sample autocorrelation function and sample partial autocorrelation function, some autocorrelation was present in the data for both seasonal and non-seasonal lags. Rather than fitting a model with lagged endogenous variables as regressors, a model known as an AR-GARCH model was selected to be fit.

This model is similar to the standard ARCH or GARCH model, but also allows for an autoregressive error structure. The standard AR-GARCH model estimated by PROC AUTOREG looks like:

$$y_t = x_t \beta + v_t$$

$$v_t = \varepsilon_t - \phi_1 v_{t-1} - \dots - \phi_m v_{t-m}$$

$$\varepsilon_t = \sqrt{h_t} e_t$$

$$h_t = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \gamma_j h_{t-j}$$

$$e_t \sim N(0,1)$$

To help determine the order of the AR process, the NLAG= and BACKSTEP options were specified in the AUTOREG procedure. The NLAG= option indicates the number of lags of the autoregressive error term to fit in the model and the BACKSTEP option removes any non-significant autoregressive coefficients. This proved to be a useful tool in obtaining a starting point for modeling the INFLATN series. The residuals from the selected model were also output to a SAS data set. Following is the corresponding SAS code.

```
proc autoreg data=meandel;
  model inflatn=/noint nlag=12 backstep;
  output out=ar_resid r=r;
  label inflatn='Monthly Inflation Rate';
  title 'determine the order of AR process';
run;
```

Figure 14. NLAG= and BACKSTEP options to help determine the order of the AR process

Figure 15 contains a portion of the corresponding listing file, and shows that the AUTOREG procedure retained lags 1,2,7,9 and 12 as significant.

Dependent Variable = INFLATN Monthly Inflation Rate			
Backward Elimination of Autoregressive Terms			
Lag	Estimate	t-Ratio	Prob
3	-0.007768	-0.1237	0.9016
10	0.013607	0.2184	0.8273
8	-0.022622	-0.3667	0.7141
4	-0.029298	-0.4908	0.6239
11	-0.045977	-0.7781	0.4372
5	-0.052031	-0.8938	0.3722
6	0.034985	0.6228	0.5339
Preliminary MSE = 6.388E-6			
Estimates of the Autoregressive Parameters			
Lag	Coefficient	Std Error	t Ratio
1	-0.36041252	0.055827	-6.456
2	-0.15937950	0.056543	-2.819
7	-0.14168419	0.052670	-2.690
9	-0.11297324	0.054475	-2.074
12	-0.10130111	0.051255	-1.976

Figure 15. Results of NLAG= and BACKSTEP options

Given this base model, the resulting residuals were then used to ascertain the presence of ARCH with the ARCHTEST option. The SAS code is contained in Figure 16 and the output is found in Figure 17.

```
proc autoreg data=ar_resid;
  model r=/noint archtest;
run;
```

Figure 16. ARCHTEST option on the residuals from AR model

Q and LM Tests for ARCH Disturbances				
Order	Q	Prob>Q	LM	Prob>LM
1	6.4281	0.0112	6.4106	0.0113
2	6.6047	0.0368	7.0145	0.0300
3	6.7810	0.0792	7.0654	0.0698
4	6.7824	0.1478	7.0681	0.1323
5	10.4629	0.0631	10.7149	0.0573
6	17.1739	0.0087	14.8509	0.0214
7	17.2561	0.0158	14.9293	0.0369
8	17.5202	0.0251	15.4556	0.0509
9	17.5328	0.0410	15.4571	0.0791
10	17.7723	0.0589	15.8415	0.1043
11	17.8615	0.0849	16.1714	0.1349
12	17.8709	0.1197	16.2727	0.1791

Figure 17. ARCHTEST results

The ARCHTEST option confirms the presence of ARCH, but does not consistently point to fitting a higher order ARCH or GARCH model. Because we have seen that the LM and Q-statistics are not especially useful in helping determine the order of the ARCH/GARCH model, we will fit several models and compare the AIC and SBC values from these competing models and select the best model from this set. Table 3 provides this summary.

Model	SBC	AIC	Non-Sign Parms
AR(1,2,7,9,12)- ARCH(1)	-2721.3	-2747.25	A(9) P=.2143
AR(1,2,7,9,12)- ARCH(2)	-2721.3	-2747.25	A(9) P=.2143 ARCH(2) P=1.0
AR(1,2,7,9,12)- GARCH(1,1)	-2721.3	-2747.25	A(9) P=.2143 GARCH1 P=1

AR(1,2,7,12)-			
ARCH(1)	-2725.6	-2747.84	
AR(1,2,7,12)-			
ARCH(2)	-2725.6	-2747.84	ARCH2 P=1
AR(1,2,7,12)-			
GARCH(1,1)	-2725.6	-2747.84	GARCH1 P=1
AR(1,2,12)-			
ARCH(1)	-2721.8	-2740.33	

Table 3. Summary of Estimated Models

Based on the models estimated, the best fitting model according to the AIC and SBC, in which all model parameter estimates were significant is the AR(1,2,7,12)-ARCH(1) model. Note that in many cases, the AIC and SBC were the same for competing models. This is because the coefficients associated with the non-significant terms in those models were very close to zero, and therefore contributed nothing to the computation of these statistics. In the interest of parsimony, the model with the fewest number of parameters and the lowest AIC/SBC value is the best.

It is interesting to note that the AR(9) coefficient, which was significant based on the OLS residuals was no longer significant when the ARCH/GARCH model was added. This upholds the argument that ignoring heteroscedasticity in data could lead to an overparameterized ARMA model.

The final model is fit as follows in Figure 18 and the residuals and conditional error variance are output using the R= and CEV= options, respectively. A portion of the output is shown in Figure 19.

```
proc autoreg data=meandel;
  model inflatn=/nlag=(1 2 7 12) garch=(q=1)
  noint;
  output out=final_r r=e cev=cev;
run;
```

Figure 18. PROC AUTOREG statments to fit best model

GARCH Estimates					
SSE	0.001961	OBS	301		
MSE	6.514E-6	UVAR	7.459E-6		
Log L	1379.919	Total Rsq	0.4442		
SBC	-2725.6	AIC	-2747.84		
Normality Test	64.1492	Prob>Chi-Sq	0.0001		
Variable	DF	B Value	Std Error	t Ratio	Approx Prob
A(1)	1	-0.407064	0.0686	-5.933	0.0001

A(2)	1	-0.147701	0.0612	-2.414	0.0158
A(7)	1	-0.148260	0.0496	-2.990	0.0028
A(12)	1	-0.088109	0.0445	-1.982	0.0475
ARCH0	1	0.000004247	5.254E-7	8.084	0.0001
ARCH1	1	0.430660	0.1113	3.871	0.0001

Figure 19. Estimates of AR(1,2,7,12)-ARCH(1) model

All of the parameter estimates are significant, the only troubling result is that the test for Normality is rejected. This could be due to the fact that the sample size is relatively small (N=301), or because the underlying distribution of the errors follows something other than a Normal distribution. To ensure that the heteroscedasticity has been accounted for, the discounted residuals are computed and run through PROC AUTOREG's ARCHTEST option. Figure 20 contains the programming statements to do this and Figure 21 holds the output. Neither the LM nor Q-statistic is significant for any of the 12 lags.

```
data final;
  set final_r;
  disc_res=e/sqrt(cev);
run;

proc autoreg data=final;
  model disc_res=/noint archtest;
run;
```

Figure 20. Compute discounted residuals and test them for any remaining heteroscedasticity

Q and LM Tests for ARCH Disturbances				
Order	Q	Prob>Q	LM	Prob>LM
1	0.0755	0.7835	0.0678	0.7946
2	0.7079	0.7019	0.6561	0.7203
3	1.0656	0.7854	1.0155	0.7975
4	1.3818	0.8473	1.3748	0.8486
5	3.6305	0.6037	3.4395	0.6326
6	9.6407	0.1406	9.3864	0.1530
7	10.0054	0.1883	9.9656	0.1905
8	10.4204	0.2367	10.8378	0.2111
9	10.6440	0.3009	10.8594	0.2855
10	10.9457	0.3618	11.0633	0.3526
11	10.9457	0.4478	11.2276	0.4244
12	11.1135	0.5192	11.2537	0.5073

Figure 21. ARCHTEST results on discounted residuals

Now that we have a reasonable model, we can compute predicted values of both the conditional mean and the conditional error variance. We can use the conditional prediction error variance to compute more accurate confidence limits about the predicted values. Note that, by default, PROC AUTOREG uses the unconditional prediction error variance when computing confidence limits for the predicted values, even when the GARCH= option is specified. This is because the conditional prediction error variance is not computable for certain types of GARCH models.

The following code in Figure 22 shows how to add 12 new observations to the INFLATN series in order to obtain forecasted values for 1 year into the future. The model is then re-run on this new set of data. Several variables are output from this model including: the predicted value of the conditional mean and its associated upper and lower confidence limits, and the conditional prediction error variance (CPEV= option). A DATA step follows which shows how to compute confidence limits based on the CPEV. These variables are then plotted.

```
data future;
  retain date '01feb95'd;
  do i = 1 to 12;
    inflatn=.;
    date = intnx('month',date,1);
    output;
  end;
  format date date7.;
run;

data total;
  set meandel future;
run;

proc autoreg data=total;
  model inflatn=/nlag=(1 2 7 12) garch=(q=1)
    noint;
  output out=final cpev=cpev p=pred_inf ucl=ucl
    lcl=lcl;
run;

data cl;
  set final;
  t=tinv(.975,301);
  ucl_c=pred_inf+t*sqrt(cpev);
  lcl_c=pred_inf-t*sqrt(cpev);
run;

symbol1 v=plus c=black;
symbol2 i=join v=none c=cyan;
symbol3 i=join v=none c=red l=2;
axis1 order=(-.015 to .030 by .015);
proc gplot data=cl;
  plot inflatn*date=1 pred_inf*date=2
    ucl*date=3 lcl*date=3/overlay
  vaxis=axis1;
  plot inflatn*date=1 pred_inf*date=2
    ucl_c*date=3 lcl_c*date=3/overlay
  vaxis=axis1;
run;
```

Figure 22. Computing the Confidence Limits

Figures 23 and 24 illustrate the difference in the confidence limits based on the unconditional prediction error variance, vs. the conditional prediction error variance. Periods of high volatility have wider confidence limits and periods of low volatility have narrower confidence limits when the conditional prediction error variance is used.




Figure 23. Confidence Limits based on unconditional variance




Figure 24. Confidence Limits based on conditional variance

ENHANCEMENTS TO ARCH/GARCH MODELING CAPABILITIES IN RELEASE 6.11

The AUTOREG procedure will support the following new ARCH/GARCH modeling features in the 6.11 release:

- Student's t distribution for the distribution of the underlying disturbances
- Nelson and Cao inequality constraints (Nelson and Cao, 1992), which are less restrictive than the non-negativity constraints
- regressor variables in the definition of the conditional error variance
- LM test for specified forms of heteroscedasticity
- supply initial values for the GARCH parameters
- estimate the start-up of the conditional error variance

- various functional forms of the GARCH-M model

CONCLUSION

This paper provided a brief introduction to the ARCH/GARCH model and showed how it can be estimated and forecast using the AUTOREG procedure. It illustrated the difficulties in determining the order of the ARCH/GARCH process and offered some practical suggestions on how to overcome this difficulty. The paper emphasized that the purpose of the family of ARCH models is not to improve the forecast of the conditional mean, but rather to use prior information to derive more accurate estimates of the forecast error variance.

The types of models described here are only a small subset of the family of ARCH models available. The AUTOREG procedure can be used to estimate the following types of ARCH models, in addition to what was described in this paper:

- integrated GARCH (I-GARCH)
- exponential GARCH (E-GARCH)
- GARCH-in-mean (GARCH-M)

Several recent papers provide a comprehensive survey of the family of ARCH models, including Bera and Higgins (1992) and Bollerslev, Chou, and Kroner (1992).

REFERENCES

- Bera, A.K and Higgins, M.L. (1992), "A Survey of ARCH Models: Properties, Estimation and Testing," Faculty Working Paper No. 92-0117, University of Illinois at Urbana-Champaign.
- Bollerslev, T. (1986), "Generalized Autoregressive Heteroskedasticity," *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T., Chou, R.Y., Kroner, K. F. (1992), "ARCH Modeling in Finance," *Journal of Econometrics*, 52, 5-59.
- Engle, R.F. (1982), "Autoregressive Conditional Heteroscedasticity Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987-1007.
- McLeod, A. I., and Li, W.K. (1983), "Diagnostic Checking ARMA Time Series Models Using Squared-

Residual Autocorrelations," *Journal of Time Series Analysis*, 4, 269-273.

Nelson, D.B., and Cao, C.Q. (1992), "Inequality Constraints in the Univariate GARCH Model," *Journal of Business & Economic Statistics*, 10, 229-235.

Weiss, A.A. (1984), "ARMA Models with ARCH Errors," *Journal of Time Series Analysis*, 5, 129-43.

SAS, SAS/ETS, and SAS/GRAPH are registered trademarks of SAS Institute, Inc. in the U.S.A. and other countries. ® indicated U.S.A. registration.